

# RELIABLE EXPLAINABILITY OF DEEP LEARNING SPATIAL-SPECTRAL CLASSIFIERS FOR IMPROVED SEMANTIC SEGMENTATION IN AUTONOMOUS DRIVING

Jon Gutiérrez-Zaballa<sup>1\*</sup>, Koldo Basterretxea<sup>1</sup>, Javier Echanobe<sup>2</sup>

Department of Electronics Technology, University of the Basque Country, Bilbao, Spain<sup>1</sup>  
Department of Electricity and Electronics, University of the Basque Country, Leioa, Spain<sup>2</sup>

## ABSTRACT

Integrating hyperspectral imagery (HSI) with deep neural networks (DNNs) can strengthen the accuracy of intelligent vision systems by combining spectral and spatial information, which is useful for tasks like semantic segmentation in autonomous driving. To advance research in such safety-critical systems, determining the precise contribution of spectral information to complex DNNs' output is needed. To address this, several saliency methods, such as class activation maps (CAM), have been proposed primarily for image classification. However, recent studies have raised concerns regarding their reliability. In this paper, we address their limitations and propose an alternative approach by leveraging the data provided by activations and weights from relevant DNN layers to better capture the relationship between input features and predictions. The study aims to assess the superior performance of HSI compared to 3-channel and single-channel DNNs. We also address the influence of spectral signature normalization for enhancing DNN robustness in real-world driving conditions.

**Index Terms**— Saliency methods, Class Activation Mapping, Hyperspectral Imaging, Semantic Segmentation

## 1. INTRODUCTION

A key challenge in computer vision is metamerism, which hinders the robustness of tristimulus based systems. Metamerism occurs when the reflectance curves of different materials match under certain conditions, like specific illumination or sensor characteristics, but differ under others [1].

Hyperspectral sensors offer a promising solution to address metamerism and increase the robustness of intelligent vision systems. These sensors capture light across narrow-band filters usually spanning visible to infrared wavelengths. Hyperspectral imaging (HSI) has proven valuable in fields such as food quality assessment, medical analysis, and remote sensing [2] and its extension to new application domains, such

as autonomous driving systems (ADS), is currently being investigated. Despite its benefits, HSI poses challenges for classification tasks due to its high dimensionality, inter-class similarity, and intra-class variability [3]. To achieve pixel-level precision in applications like semantic or instance segmentation, combining spectral and spatial information is crucial. One approach to achieving this is by using fully convolutional networks (FCNs), a type of deep neural network (DNN) with an encoder-decoder structure.

To comprehend the contribution of spectral information beyond the numerical evaluations provided by ablation experiments, DNNs must be interpretable, meaning we should be able to interpret their behavior in the decision-making processes. Explainable AI, which focuses on the interpretability of neural networks, is a hot topic nowadays [4]. In the intelligent vision area, the most common methods for identifying which image features are most relevant for predictions are the saliency methods, including layer-wise relevance propagation (LRP), gradient-based techniques, class activation mapping (CAM), and their variants. However, despite recent advancements, many of these methods exhibit notable limitations. As noted by [5, 6, 7, 8], most saliency methods for image classification are designed to produce visually appealing representations rather than accurately reflecting the underlying prediction processes.

In this article, in Section 2, we first review the desirable properties that saliency methods should have to be considered reliable, as discussed in the literature. In Section 3, we extend the CAM conservativeness property –originally defined for classification networks– to semantic segmentation networks. Next, in Section 4, we examine the essential characteristics a CAM-based saliency method for segmentation must have to satisfy this property, demonstrating that many such methods do not adequately fulfill it. To overcome these limitations when interpreting the inference process on FCNs, in Section 5 we propose a method that leverages activations and weights from specific, relevant model layers to more reliably explain the relationship of input information with predictions. We apply this method to evaluate a segmentation U-Net trained with images from a 25-band hyperspectral snapshot camera in real-world driving scenarios, assessing how spectral information contributes to the enhanced performance of HSI-based clas-

\*This work was partially supported by the University of the Basque Country (UPV-EHU) under grant GIU21/007, by the Basque Government under grants PRE\_2023\_2.0148 and KK-2023/00090, and by the Spanish Ministry of Science and Innovation under grant PID2020-115375RB-I00.

sification models. Finally, Section 6 provides a discussion of our findings and their implications. It also outlines directions for future work.

## 2. BACKGROUND

### 2.1. Invariance Tests for Saliency Methods

The reliability of saliency methods has recently been questioned, prompting efforts to identify the invariants that saliency methods should satisfy to be considered trustworthy. As highlighted by the authors in [5], some widely used techniques are independent of model parameters and training data, thus failing to establish valid relationships between inputs and outputs present in the data.

In this context, [5] conducts two sanity checks on various saliency methods: one based on model parameter randomization (comparison against a randomly initialized untrained network) and the other on data randomization (comparison against a trained network but with randomly permuted labels). GradCAM passes these sanity checks, while Guided Back-Prop and GuidedGradCAM are invariant to parameters in higher layers, hence failing. However, it is important to note that when dealing with overparameterized, non-optimized models (unpruned/unquantized models), certain layers in the network may contribute minimally to the output, leading to some level of invariability to parameter perturbations. Similarly, the authors of [7] show, through adversarial examples, that LRP fails to explain the DNN’s decision-making process for original images or adversarial examples, posing a further challenge for reliable explainability methods.

The authors of [6] propose the requirement of input invariance to ensure a reliable interpretation of the input’s contribution to the model predictions. This principle dictates that a saliency method should reflect the model’s sensitivity to transformations of the inputs. Finally, in [8], the authors introduce the concept of conservativeness in CAM. Conservativeness implies that the sum of the contributions in the mapping aligns with the prediction scores. The authors show that the original CAM [9] satisfies this property, while GradCAM [10] does not, leading to the proposal of a variant method called Extended-CAM.

### 2.2. Basics of CAM-based Studies

All methods based on class activation mappings derive from the original CAM [9]. The goal is to investigate the contribution of the spatial component  $(i, j)$  to the final prediction  $y^c$  for an image classified as belonging to class  $c$ . As previously mentioned, the sum of these contributions ( $L_{ij}^c$ ) should match  $y^c$  to satisfy the conservativeness property (Eq. 1).

$$\sum_{i,j} L_{i,j}^c = y^c \quad (1)$$

However, the original CAM is limited by the assumption that the last layer of the feature extractor consists of a global average pooling layer followed by a fully-connected layer, which restricts its applicability to modern DNNs.

To extend the use of CAM to more generic classification networks, GradCAM (Eq. 2) was introduced in [10]:

$$L_{i,j}^c = ReLU\left(\sum_k A_{ijk} \alpha_k^c\right) \quad (2)$$

where  $A_{ijk}$  is the activation of a certain layer and  $\alpha_k^c$  is the linear multiplying coefficient. Alternative methods based on CAM propose different ways to calculate these  $\alpha_k^c$  coefficients. In the original GradCAM, these coefficients are determined through a global average pooling operation (Eq. 3):

$$\alpha_k^c = \frac{1}{Z} \sum_u \sum_v \frac{\partial y^c}{\partial A_{u,v}^k} \quad (3)$$

The corresponding approach for segmentation tasks, SegGradCAM, was proposed in [11] and is detailed in (Eq. 4).

$$L_{i,j}^c = ReLU\left(\sum_k A_{ijk} \frac{1}{N} \sum_{(r,s) \in M} \frac{\partial \sum_{(r,s) \in M} y_{rs}^c}{\partial A_{uv}^k}\right) \quad (4)$$

where  $M$  is the set of pixels in the region of interest (RoI), and  $y_{rs}^c$  is the non-probabilistic output (logit) for the pixel at  $(r, s)$  belonging to class  $c$ . When  $M$  contains only one pixel, SegGradCAM (Eq. 4) simplifies to GradCAM (Eq. 2), even though they describe different input/output connections.

## 3. EXTENDING CONSERVATIVENESS TO SEMANTIC SEGMENTATION

Conservativeness ensures the precise quantitative explanation of prediction scores by providing absolute pixelwise contributions without redundancy or deficiency [8]. This property is crucial for visualization methods in neural networks focused on semantic segmentation, as it allows for analyzing the input contributions to a single output pixel, specific pixels in a RoI, or every pixel associated with a certain label. If this condition is not met, the validity of comparing CAMs from different layers or DNNs becomes questionable.

In [8], conservativeness is proposed for image classification models. We extend this property to semantic segmentation models in Eq. 5 which indicates that a pixel activation mapping  $L_{i,j,r,s}^c$  for class  $c$  and output location  $(r, s)$  is considered conservative if and only if the spatial sum over  $(i, j)$  equals the target score  $y_{r,s}^c$  for every pixel.

$$\sum_{i,j} L_{i,j,r,s}^c = y_{r,s}^c \quad (5)$$

Class-level conservativeness is defined in Eq. 6 and depends on pixel-level conservativeness and vice versa.

$$\sum_{r,s \in M} (y_{r,s}^c \cdot m_{r,s}) = \sum_{r,s \in M} [(\sum_{i,j} L_{i,j,r,s}^c) \cdot m_{r,s}] = Y^c \quad (6)$$

where  $y_{r,s}^c$  is the logit for the pixel at position  $(r, s)$  belonging to class  $c$ ,  $m(r, s)$  is a 0-valued matrix with the same size as the output, containing a 1 at  $(r, s)$  and  $Y^c$  is the logit map for class  $c$ . Finally, the segmentation class activation mapping is defined as in Eq. 7.

$$L_{i,j}^c = \sum_{r,s \in M} L_{i,j,r,s}^c \quad (7)$$

The interpretation of Eqs. 5, 6 and 7 is illustrated in Fig.

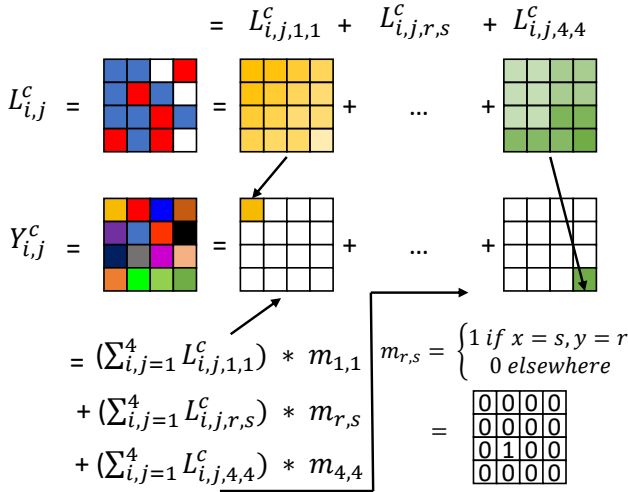


Fig. 1: Conservativeness property in semantic segmentation.

#### 4. CONSERVATIVENESS ASSESSMENT IN GRADCAM- AND SEGGRADCAM-BASED METHODS

As previously discussed in Section 2, many saliency methods do not adhere to fundamental invariants. Hereafter, we show that several widely used CAM-based methods, included some designed for classification, also fail to meet the conservativeness principle. The selection of these specific methods is motivated by their prominence in the literature and their unique approaches for calculating  $\alpha_k^c$  coefficients (Eq. 3).

To illustrate the non-conservativeness of gradient-based saliency methods, we just can rewrite Equation 1 in the form of GradCAM (Eq. 2), which is known not to meet the conservativeness criterion due to the spatial-averaging global pooling operation performed (Eq. 3), as noted in [8].

SegGradCAM (see Subsection 2.2) modifies Eq. 4 into Eq. 8 by applying the commutative property between partial derivatives and summation, while omitting the ReLU operation for simplicity.

$$L_{i,j}^c = \sum_{(r,s) \in M} L_{i,j,r,s}^c = \sum_{(r,s) \in M} \sum_k A_{ijk} \left( \frac{1}{N} \sum_{u,v} \frac{\partial y_{rs}^c}{\partial A_{uvk}} \right) \quad (8)$$

When  $M$  contains multiple pixels, each pixel has a specific logit activation map  $L_{i,j,r,s}^c$ , which is added to obtain the logit map of class  $c$ , as illustrated in Fig. 1. For a single pixel in  $M$ , the equations for SegGradCAM and GradCAM are essentially identical, making the demonstration of non-conservativeness straightforward.


The use of global average pooling over the gradient in calculating  $\alpha_k^c$  for semantic segmentation introduces unwanted effects. For instance, using an image from HSI-Drive [12], the CAM for the Road Marks class (Fig. 2b) shows nearly every pixel contributing, which is counterintuitive. To address this, we can eliminate the average from  $\alpha_{ijk}^c$  calculation, giving spatial dimensions to it. The equation for SegGradCAM with the corrected  $\alpha_{ijk}^c$  is shown in Eq. 9.

$$L_{i,j}^c = \sum_{(r,s) \in M} L_{i,j,r,s}^c = \sum_{(r,s) \in M} \left( \sum_k A_{ijk} \odot \frac{\partial y_{rs}^c}{\partial A_{ijk}} \right) \quad (9)$$

This correction yields a more meaningful class activation mapping for the Road Marks class, as shown in Fig. 2c. However, a more reasonable appearance does not ensure adherence to the conservativeness principle. Our evaluation demonstrates that Eq. 9 can also be derived by extending Extended-CAM [8] formulation, which satisfies conservativeness principle for classification, to the segmentation task.

The authors of ExtendedCAM suggest clipping negative values of the class activation mapping using ReLU functions, claiming it enhances visualization accuracy. However, this practice violates the conservativeness principle and risks losing critical information regarding regions contributing negatively to the classification. Furthermore, for ExtendedCAM and its extension to segmentation to maintain conservativeness, nonlinear layers of the network must be approximable by the first term of a Taylor decomposition. This is a condition not generally applicable to segmentation DNNs employing standard data processing structures, including Max-Pooling and Concatenation layers like those in U-Net.

The SegXResCAM method proposed in [13] employs pooling/unpooling operations on gradients to regulate their granularity. Nevertheless, similar to the average pooling



**Fig. 2:** Effect of the SegGradCAM gradient weighting modification on *Conv2D\_1* activation from HSI & PN model. Image 1111\_576 from HSI-Drive v2.0 [12]. a) Pseudocolor version. b) Spatial average. c) Non-spatial average.

in GradCAM, this spatial processing compromises conservativeness. The classification non-conservativeness of SegXResCAM is evidenced in Eq. 10.

$$\sum_{i,j} L_{i,j}^c = \sum_{i,j} ReLU\left(\sum_k Up\left[Pool\left[\frac{\partial Y^c}{\partial A_{ijk}}\right]\right] \odot A_{ijk}\right) \neq y^c \quad (10)$$

In what follows, we also evaluate RelevanceCAM [14] and ScoreCAM [15], which are not gradient-based but also fail to satisfy conservativeness.

RelevanceCAM (Eq. 11) does not meet conservativeness because it derives weights from a global average pooling operation over the relevance maps  $R^{(c,i)}(x, y)$ .

$$\sum_{i,j} L_{i,j}^c = \sum_{i,j} \sum_k \left(\sum_{x,y} R_k^{(c,i)}(x, y)\right) A_k^c \neq y^c \quad (11)$$

ScoreCAM (Eq. 12) defines the weighting of activation maps  $C(A_l^k)$  through a channel-wise increase of confidence score based on perturbation analysis. Its definition, differing from the network's inference process, results in non-conservativeness.

$$\sum_{i,j} L_{i,j}^c = \sum_{i,j} ReLU\left(\sum_k C(A_l^k) A_l^k\right) \neq y^c \quad (12)$$

The absence of conservativeness across these methods impedes performing meaningful comparative analysis, particularly when evaluating CAMs across multiple models in our study. While these techniques can still be valuable to understand model behavior, particularly for simpler classification tasks, this article emphasizes that saliency methods should focus on representing reality as accurately as possible and produce insightful heat maps rather than fancy, superficial visualizations.


## 5. ANALYZING THE CONTRIBUTION OF SPECTRAL INFORMATION IN HSI SEGMENTATION DNN MODELS

### 5.1. A DNN for Scene Understanding Using HSI

HSI-Drive [12] is a comprehensive dataset designed specifically for researching HSI technology in developing more accurate and robust ADS, particularly under adverse environmental conditions. We trained and optimized various U-Net-based machine learning models on HSI-Drive using different classification targets to create image segmentation systems of varying complexity. To assess the contribution of spectral information to the inference processes of these models, we evaluated the performance of the same model architecture using 1, 3, and all the 25 spectral channels provided by an Imec 5x5 mosaic-type near-infrared (NIR) HSI sensor. Fig. 3 shows an example of the 5-class segmentation experiment comprising Road (tarmac), Road Marks, Vegetation, Sky, and "Other" classes. For more in-depth information regarding the DNN architecture and the dataset used in this article, the reader is referred to [12, 16].

HSI-Drive was recorded with a snapshot camera equipped with a multispectral filter array (MSFA) based on Fabry-Perot interferometers. This technology, when used with no rejection filters, as was the case, can lead to some input channels containing mixed contributions from distinct wavelengths [12]. For this study, 1-channel images consist of the MSFA central band, which exhibits a spectral response centered around 731.883 nm. 3-channel images, which were generated as an alternative to RGB imaging (no visible RGB bands can be extracted from these images), consist of the three most informative spectral bands, which were extracted using the orthogonal space projection method as explained in [17]. Specifically, these bands include: 9 (with peak at 770.576 nm), 22 (featuring a primary peak at 858.948 nm and a secondary peak at 577.147 nm), and 25 (having a primary peak at 944.485 nm and a secondary peak at 657.995 nm).

For the 25-band images, we applied per-pixel normalization (PN) in the preprocessing stage, which involves dividing each of the 25 spectral components by the total sum of all.



**Fig. 3:** Inference result on image 3112\_104 from HSI-Drive [12].

By applying PN, we prioritize the shape of the spectral signatures over variations in mean reflectance values. Although mean reflectance data can be very useful for class discrimination under constant and controlled illumination conditions, it tends to perform poorly under varying lighting and severely shadowed scenes, as we could observe in the analysis of the HSI-Drive dataset. Hence, even though this article is not focused on pixel normalization techniques, we observed that while applying PN to the 25-band images provides the best classification results, networks trained with 3-channel and 1-channel images without PN result in generally more accurate predictions than for their normalized counterparts.

**Table 1:** Test IoU metrics with the three models under study.

	Road	Marks	Veg.	Sky	Other
<b>1-ch.</b>	97.75	91.04	89.34	91.94	70.29
<b>3-ch.</b>	98.48	93.24	91.90	92.86	77.23
<b>25-ch. (PN)</b>	98.11	89.27	94.12	92.00	80.94

## 5.2. Activation and Weight Study for DNN Explainability

Since none of the saliency methods described in Section 4 provide reliable information for comparing the three models under study (1, 3 and 25 spectral channels), we propose to focus the explainability analysis on features that directly contribute to the computation of the final prediction values in the model.

Firstly, we examined the activations before the final convolution layer ( $conv2D_{21}$ ) and the weights and biases of the last convolutional layer ( $conv2D_{22}$ ). The output of  $conv2D_{21}$  consists of 32 non-negative feature channels with the same size as the input. Each of the 5 output class nodes has a set of 32 weights and 32 biases. The first analysis is aimed at determining the amount of relevant feature maps per class, identifying any features that contribute similarly across multiple classes, and finding features with significant positive or negative activations for specific classes.

Due to skip connections in the U-net model [16], the second convolutional block ( $conv2D_{11}$ ) is directly linked to the final convolutional block. This helps to preserve spatial information lost during subsequent encoding layers, thus providing

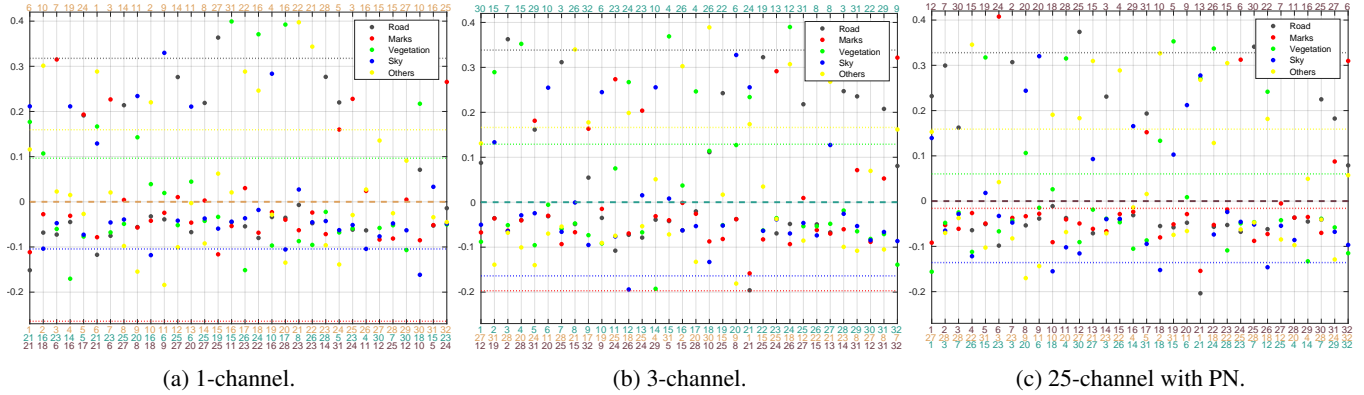
particularly meaningful information for this study. In consequence, secondly we investigated the activations at this initial layer, which usually extract relevant features for edge detection, depending on the number of spectral channels provided at the input.

### 5.2.1. Activation of $conv2D_{21}$

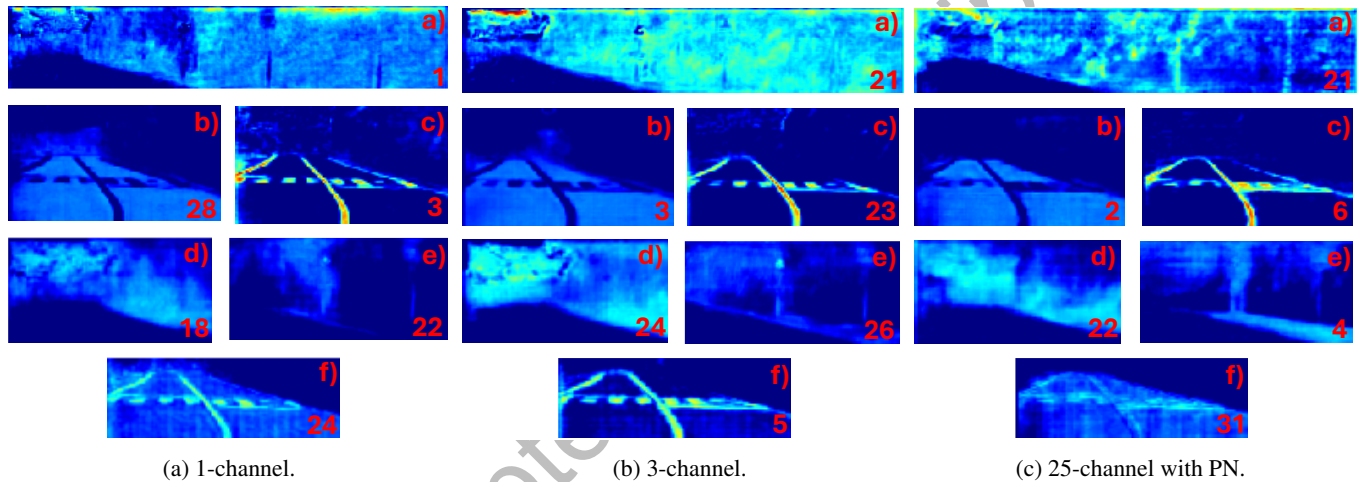
For a comparative visual analysis, we have selected and cropped six activations from this layer, shown in Fig. 5. To ensure a fair comparison across the three models, we focused on the three most correlated activations. This election was made by calculating the correlations between the weight vectors of each channel, as shown on the X-axis of Fig. 4. For example, channel 21 of the HSI model shows the highest correlation with channel 1 of the 1-channel model and channel 21 of the 3-channel model. Across all models, weights associated with Vegetation (green dot), Sky (blue dot), and Others (yellow dot) hover around 0.2, while weights linked to Road (black dot) and Road Marks (red dot) are approximately -0.15.

Row *a* of Fig. 5 compares the outputs of three correlated channels that are typically positively activated for Vegetation, Sky, and Other classes, while being negatively activated for Road and Road Marks classes. It can be observed that the activations in the 25-channel model are less homogeneous in highly diverse areas of the image, indicating an effective utilization of the spectral information. The advantages of combining HSI information and PN, which aim to homogenize regions partially in sun and shadow, are evident in Figs. 3a to 3c and row *d* from Fig. 5.

In row *b* of Fig. 5, which shows strong activation for the Road class, we observe certain regions of Vegetation that are erroneously activated in both the 1-channel and, to a lesser extent, 3-channel models. This highlights a significant drawback of PN: the lack of edge definition, which leads to poorer segmentation of the Road Marks class (Table 1). An analysis of the channels that are highly activated for the Road Marks class (row *c* from Fig. 5) and the channels activated for both the Road and Road Marks classes (row *f* from Fig. 5) reveals this issue further. While the borders are less sharp for the 25-band model, it is noteworthy that there are no erro-



**Fig. 4:** Weight (dots) and bias (dashed line) values of the  $1 \times 1$   $conv2D_{22}$  convolution layer for each class in the three models. Lines below the graph indicate the most correlated weights from other models, while the line above indicates the most correlated weights from the same model: orange (1-channel), turquoise (3-channel), and brown (25-channel with PN).



**Fig. 5:** Selected outputs/inputs from block  $conv2D_{21/22}$  of 1-channel (left), 3-channel (center) and 25-channel with PN (right) models. Numbers indicate the activation channel. Letters indicate highly correlated channels as shown in Fig. 4. a) Common activation for Veg., Sky and Others. b) Strong activation for Road. c) Strong activation for Marks. d) Homogeneity of Veg. under strong light contrast. e) Strong activation for Others. f) Common activation for Road and Marks.

neous activations outside the pavement in row *c* from Fig. 5, a problem observed in the 1-channel and 3-channel models. Another significant finding in Fig. 4 is that the number of channels positively activated for Road Marks (with a weight of more than 0.1) is lower for the the 25-channel model with PN, which is compatible with its relative lack of sharpness in the edge detection.

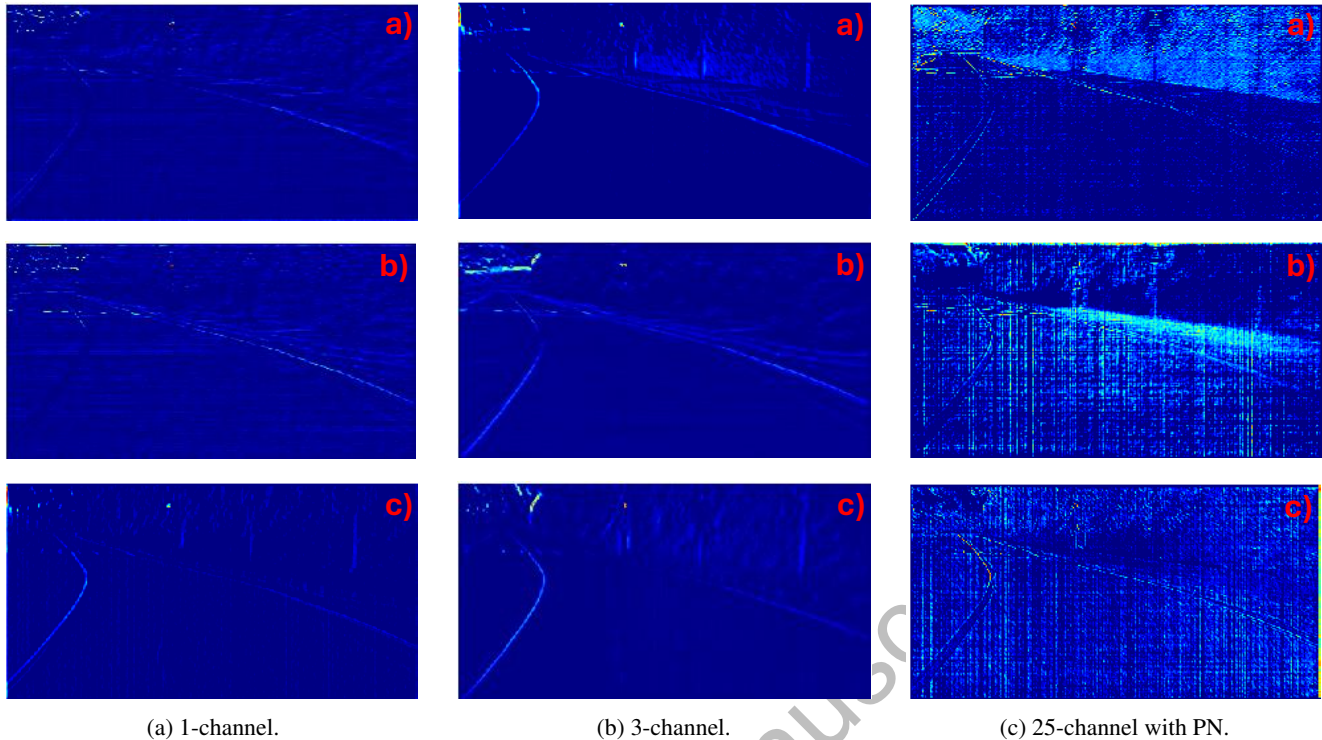
Finally, regarding the separation between the sidewalk (Others) and Vegetation, which is incorrectly classified in the 1-channel case and poorly in the 3-channel case (with the sidewalk classified as Vegetation, as in Fig. 3a and Fig. 3b), we see that even in the strongest activation for the Others class, there are no indications of correct sidewalk identification (row *e* from Fig. 5). Nevertheless, this demonstrates how normalization diminishes the impact of illumination dif-

ferences in segmentation. In this final activation, the traffic light is hardly activated, contrasting with its activation in the 1-channel and especially the 3-channel cases.

### 5.2.2. Activation of $conv2D_1$

In this case, we used a different criterion to select the channels to display, since the convolutional kernels from layer  $conv2D_2$  are  $3 \times 3$  instead of  $1 \times 1$ . We examined the 32 activations from the three models and identified the three most significant ones that exhibit the greatest similarity as shown in Fig. 6.

Due to PN normalization, as seen in Fig. 6c, the activations are noisier and more heterogeneous, leading to defective edge detection of road lines (such as curvy lines and ze-



**Fig. 6:** Selected outputs from block *conv2D\_1* of 1-channel (left), 3-channel (center) and 25-channel with PN (right) models. a) Border detection in Vegetation. b) Border detection in Zebra crossing. c) Border detection in curvy Road Mark.

bra crossings) compared to the 1-channel case, Fig. 6a, and the 3-channel case, Fig. 6b. However, this also benefits the 25-channel PN model, enabling better separation between the Vegetation and Others classes from the early layers—an effect that is imperceptible in the 1-channel and 3-channel cases. Additionally, this approach ensures more homogeneous activation in vegetation areas that are under both sunlight and shadow, in contrast to the edge detection presented in Fig. 6b for the 3-channel case.

## 6. DISCUSSION AND FUTURE WORK

As shown in this article, saliency methods show significant limitations for the explainability of DNNs for image segmentation tasks. We reviewed various tests and invariants proposed in the literature, including network and data randomization, class activation mapping conservativeness, and adversarial examples, to assess the reliability of these methods.

In particular, by extending the definition of conservativeness from image classification to image segmentation, we show that CAM-based methods, which differ in their approaches to obtaining  $\alpha_k^c$ , do not meet the conservativeness criteria. This insight suggests that the utility of these methods for precise model interpretation is questionable. Furthermore, we observed that most proposed saliency methods prioritize visually appealing representations over accurate depictions of

the actual inference processes in the models.

Given these limitations, and with the aim of exploring the advantages of using HSI in segmentation tasks, we propose to perform explainability analysis in segmentation FCNs by analyzing those feature extraction layers that directly contribute to the output predictions. Although limited, we think this approach assures a higher level of reliability in the analysis. Specifically, we analyzed activations from the penultimate layer to identify the most strongly triggered activations for each class and from the second layer to evaluate the edge detection process in segmentation tasks. Our findings indicate that models utilizing a greater number of spectral channels can enhance segmentation performance for classes with high intraclass variability.

This 25-channel model was trained using PN, which favors information provided by spectral signatures rather than general reflectance contrasts. However, PN can lead to less accurate edge detection, particularly for classes that more strongly rely on reflectance differences. This is particularly evident in the analysis of the second-layer activations, where, unlike in the 1-channel and 3-channel models, the PN method activates many regions that do not correspond to theoretical edges. However, the model demonstrates greater robustness to changes in illumination, as evidenced by consistent segmentation of regions under both sunlight and shadow.

These insights highlight the need for ongoing research

to develop innovative strategies that effectively leverage the strengths of spectral information while ensuring robust model performance across varying lighting conditions. Moreover, there is a critical need for better interpretability models specifically tailored for semantic segmentation networks to enhance understanding and trust in these systems. Future work should focus on creating models that balance the richness of data provided by HSI with the ability to generalize effectively, thereby improving the accuracy and reliability of segmentation outcomes in practical applications, particularly for autonomous driving systems.

## 7. REFERENCES

- [1] William David Wright, "The measurement of colour," 1944.
- [2] Anuja Bhargava, Ashish Sachdeva, Kulbhushan Sharma, Mohammed H. Alsharif, and Peer-aping Uthansakul et al., "Hyperspectral imaging and its applications: A review," *Heliyon*, p. e33208, 2024.
- [3] Mercedes Eugenia Paoletti, Juan Mario Haut, Javier Plaza, and Antonio Plaza, "Deep learning classifiers for hyperspectral imaging: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 158, pp. 279–317, 2019.
- [4] Sajid Ali, Tamer Abuhmed, Shaker El-Sappagh, Khan Muhammad, and Jose M et al. Alonso-Moral, "Explainable artificial intelligence (xai): What we know and what is left to attain trustworthy artificial intelligence," *Information fusion*, vol. 99, pp. 101805, 2023.
- [5] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, and Moritz et al. Hardt, "Sanity checks for saliency maps," *Advances in neural information processing systems*, vol. 31, 2018.
- [6] Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, and Kristof T et al. Schütt, "The (un) reliability of saliency methods," *Explainable AI: Interpreting, explaining and visualizing deep learning*, pp. 267–280, 2019.
- [7] Tamara R Dieter and Horst Zisgen, "Evaluation of the explanatory power of layer-wise relevance propagation using adversarial examples," *Neural Processing Letters*, vol. 55, no. 7, pp. 8531–8550, 2023.
- [8] Bum Jun Kim, Gyogwon Koo, Hyeyeon Choi, and Sang Woo Kim, "Extending class activation mapping using gaussian receptive field," *Computer Vision and Image Understanding*, vol. 231, pp. 103663, 2023.
- [9] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.
- [10] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, and Devi et al. Parikh, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [11] Kira Vinogradova, Alexandr Dibrov, and Gene Myers, "Towards interpretable semantic segmentation via gradient-weighted class activation mapping (student abstract)," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 10, pp. 13943–13944, Apr. 2020.
- [12] Jon Gutiérrez-Zaballa, Koldo Basterretxea, Javier Echanobe, M. Victoria Martínez, and Unai Martínez-Corral, "Hsi-drive v2.0: More data for new challenges in scene understanding for autonomous driving," in *2023 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2023, pp. 207–214.
- [13] Syed Nouman Hasany, Caroline Petitjean, and Fabrice Mériaudeau, "Seg-xres-cam: Explaining spatially local regions in image segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 3733–3738.
- [14] Jeong Ryong Lee, Sewon Kim, Inyong Park, Taejoon Eo, and Dosik Hwang, "Relevance-cam: Your model already knows where to look," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 14944–14953.
- [15] Haofan Wang, Mengnan Du, Fan Yang, and Zijian Zhang, "Score-cam: Improved visual explanations via score-weighted class activation mapping," *CoRR*, vol. abs/1910.01279, 2019.
- [16] Jon Gutiérrez-Zaballa, Koldo Basterretxea, and Javier Echanobe, "Evaluating single event upsets in deep neural networks for semantic segmentation: An embedded system perspective," *Journal of Systems Architecture*, vol. 154, pp. 103242, 2024.
- [17] Jon Gutiérrez-Zaballa, Koldo Basterretxea, Javier Echanobe, M. Victoria Martínez, and Unai Martínez-Corral et al., "On-chip hyperspectral image segmentation with fully convolutional networks for scene understanding in autonomous driving," *Journal of Systems Architecture*, vol. 139, pp. 102878, 2023.